US CMS L1 Trigger Hardware R&D

Thomas A. Gorski, Wesley H. Smith, U. Wisconsin - Madison

Trigger Technical Review

Fermilab, Aug. 28-29, 2017

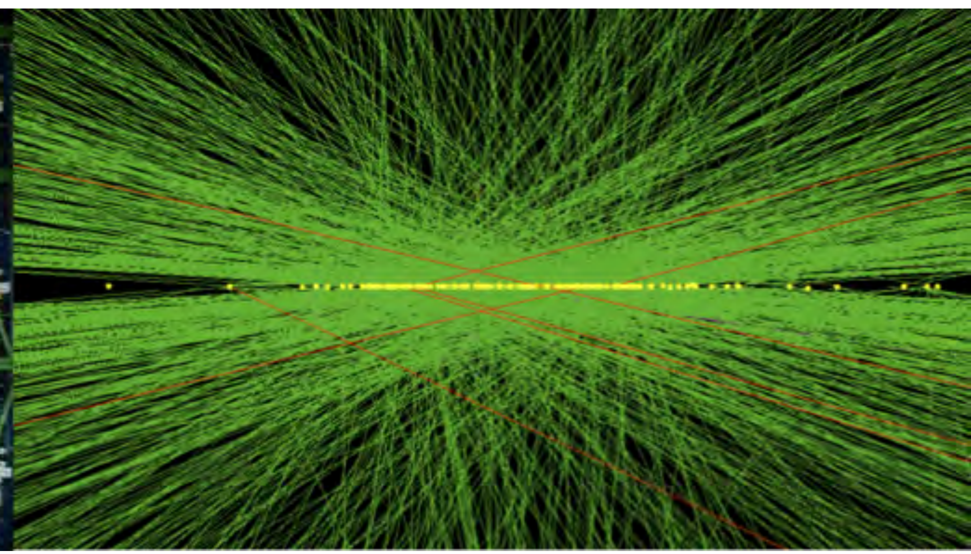

Outline

- Barrel Calorimeter Design Overview Reminder
- Barrel Calorimeter Trigger Details
- Barrel Calorimeter Trigger Layout
- Calorimeter and Correlator Trigger Demonstrator
- Calorimeter and Correlator Trigger R&D Plan
- Summary

BCal Trigger Design Reminder

- Goal: Dimension a calorimeter trigger architecture using existing or under-development technologies.
  - FPGAs: Xilinx Ultrascale and Ultrascale+ families.
  - Optics: Samtec Firefly Modules – 100Mbps to **16 Gbps**.
    - Either 12 transmitters or 12 receivers per module.
    - 14.1 Gbps modules already available, **16 Gbps under development**.
    - Each link allows up to **352bits/BX** of payload, assuming 16 Gbps line rates, 64b66b encoding and 32bits/packet reserved for protocol.
  - Build upon Phase-1 experience with hardware, firmware, software
- Close ties between algorithm development, simulation studies, firmware and software development and design engineering to provide a hardware platform for High-Luminosity LHC physics.
  - Exploit new High Level Synthesis (HLS) tools (algorithm talk later)

Inputs and Outputs

- Inputs:
  - ECAL **crystal** level information (5x5 crystals per tower) - assuming **16bits/crystal** or 400bits for one 1x1 region.
  - HCAL **tower** level information – assuming **16bits/tower**.
  - Refer to Barrel Calorimeter talks for more information.

1x1 ECAL region

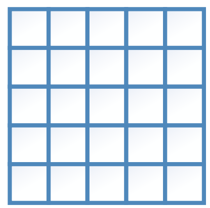

at the trigger

1x1 HCAL region

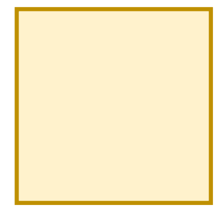

at the trigger

- Outputs:
  - Cluster objects will be sent to the Correlator.
  - Triggerable objects (standalone calorimeters triggers) will be sent to the Global Trigger.

Context Diagram

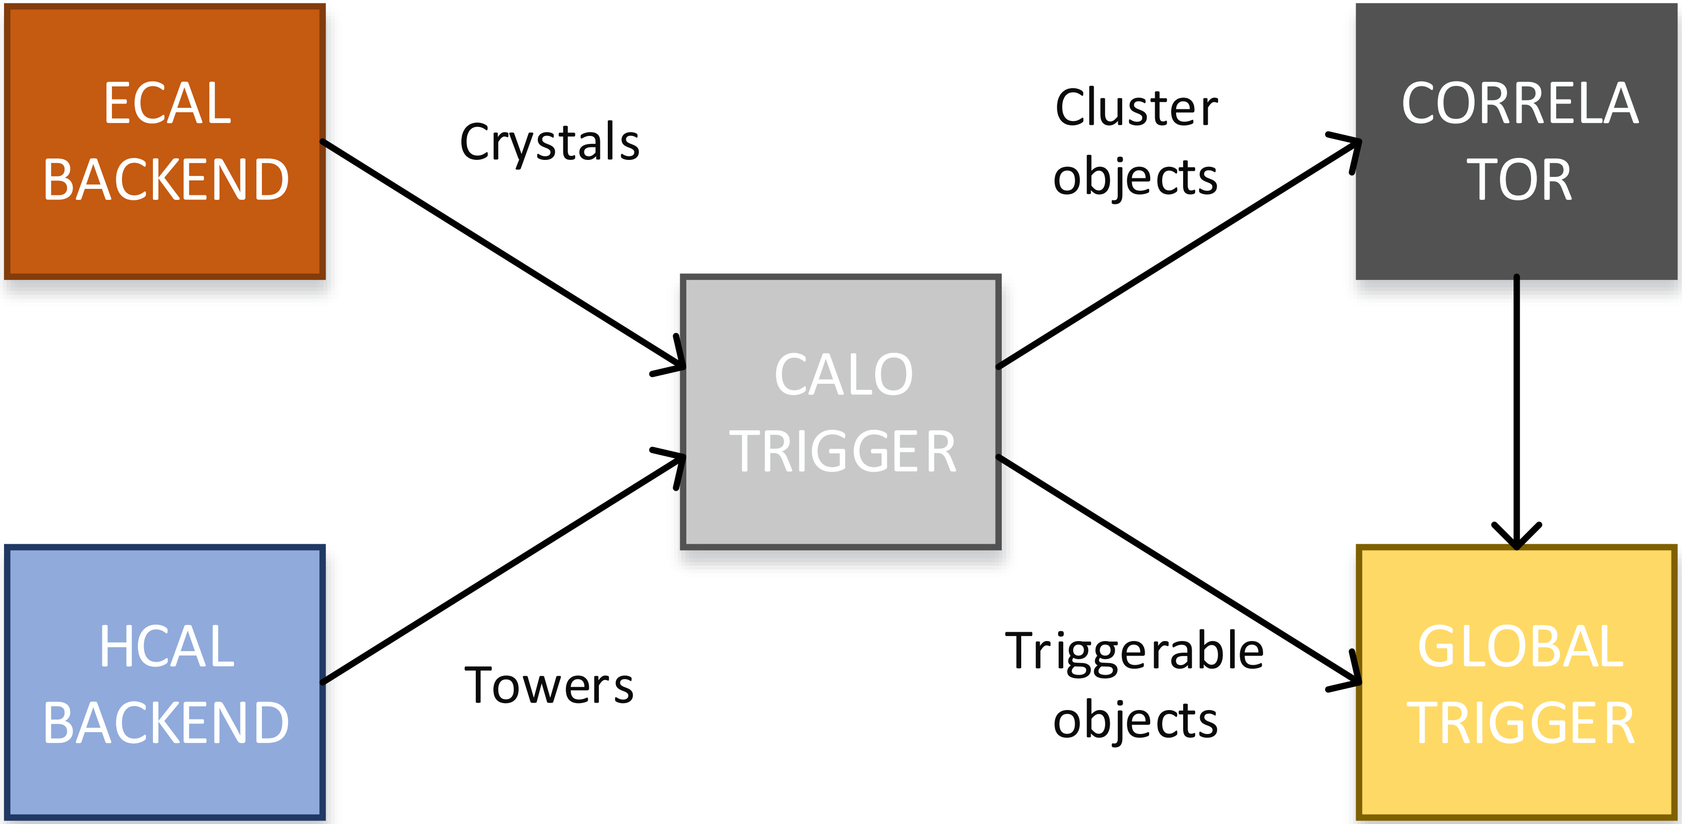

Architecture

- Use a tiled multi-layer architecture where:
  - Layer-1 partitions the detector and forms regional clusters.
  - Layer-2 stitches neighbouring clusters and forms detector-wide triggerable objects (e.g. MET).
  - Possibility to expand by adding additional layers or more cards to a certain layer.

- Designed based on the Xilinx C2104 package:
  - Package supports 104 links, 96 targeted for optical I/O.
  - Remaining 8 links reserved for DAQ, control, etc.

Calorimeter outputs detail

- ECAL:
  - Back-end divided in **3η x 4ϕ** and **2η x 4ϕ** regions, each sending **15 and 10 fibers** respectively at 16Gbps with crystal level information at **16bit/crystal**.
  - Total of 216 regions, each processed by a single FPGA.
  - Each ECAL back-end card will have 2 FPGAs, total of 108 cards.
- HCAL:
  - Back-end divided in **16η x 4ϕ** regions and tower level energies are sent out with 16Gbps links at 16bit/tower.
  - Total of 36 regions, each processed by a single FPGA.
  - Fiber count will depend on how the trigger is partitioned.
    - Will match ECAL regions – ~1 fiber for a **3η x 4ϕ** region..

ECAL and HCAL BE regions

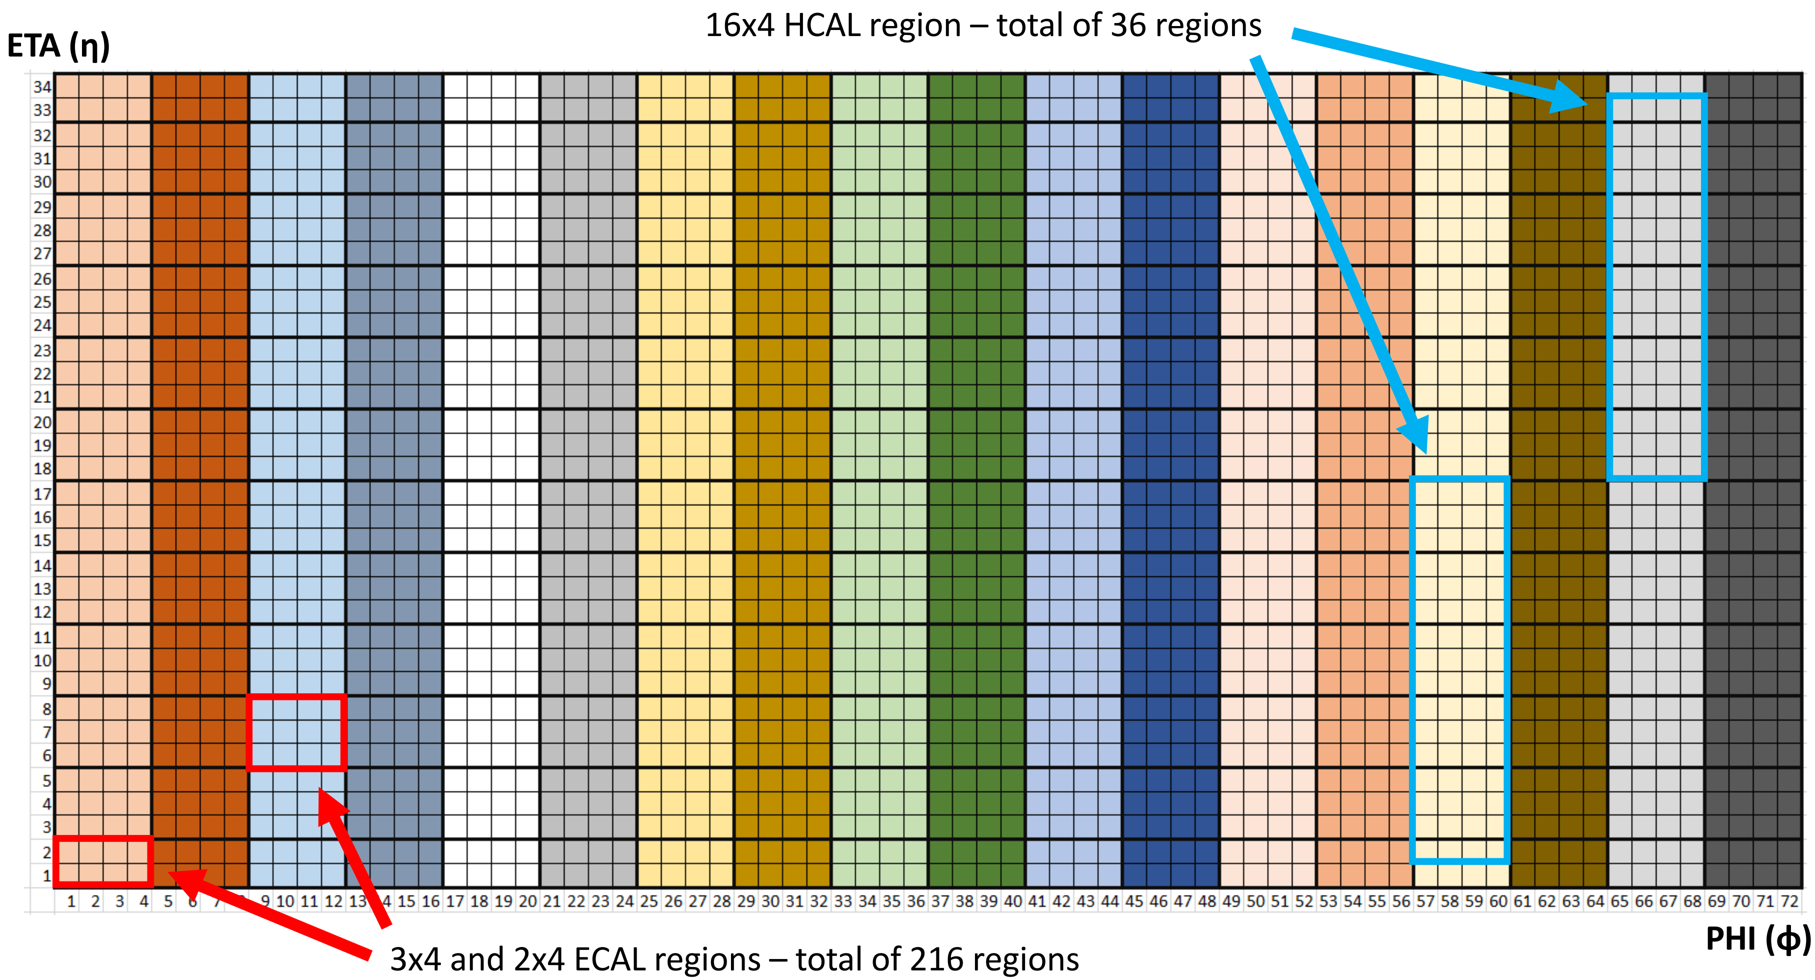

System Layout

- **Boards with 96 optical links available for data reception and transmission.**
  - **4 Additional links for DAQ readout**
- Layer-1 partitions detector in **17η x 4ϕ** regions – total of 36 regions.
  - ECAL inputs: 5x (**3η x 4ϕ)** and 1x (**2η x 4ϕ)** regions - 85 ECAL fibers.
  - HCAL inputs: 1x (**16η x 4ϕ)** region - 4 HCAL fibers.
  - Outputs: 6 fibers per region with regional clusters and metadata (2.1kbits/BX).
- Layer-2 partitions the detector **34η x 24ϕ** regions – total of 3 regions. Data duplication between Layer-1 regions required.
  - Inputs from Layer-1: 6 fibers x 12 **(34η x 24ϕ)** regions - 72 Layer-1 fibers.
  - From neighbours: 6 fibers x 4 **(34η x 24ϕ)** regions – 24 Layer-1 fibers.
  - 288 outputs available for clusters for the correlator and standalone trigger objects for the Global Trigger.
- A total of 36 layer-1 and 3 layer-2 Boards are required: **39 Cards**.
  - A total of **288 fibers** are required between layers.

System Layout Geometry (1)

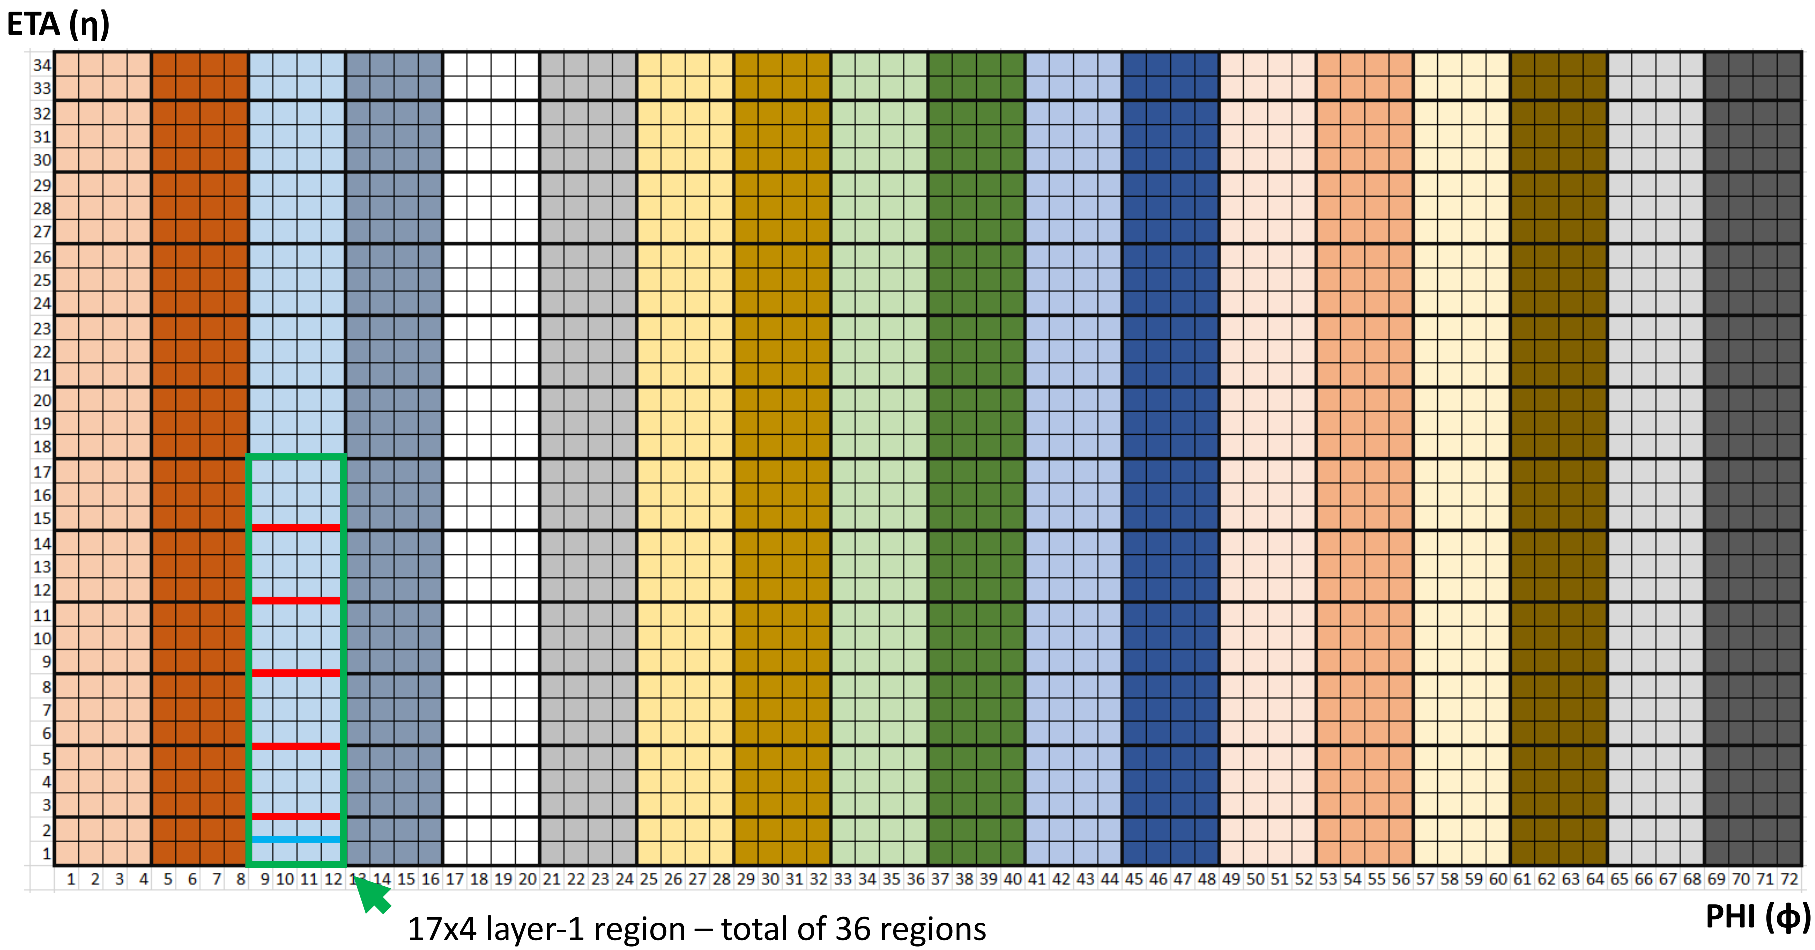

System Layout Geometry (2)

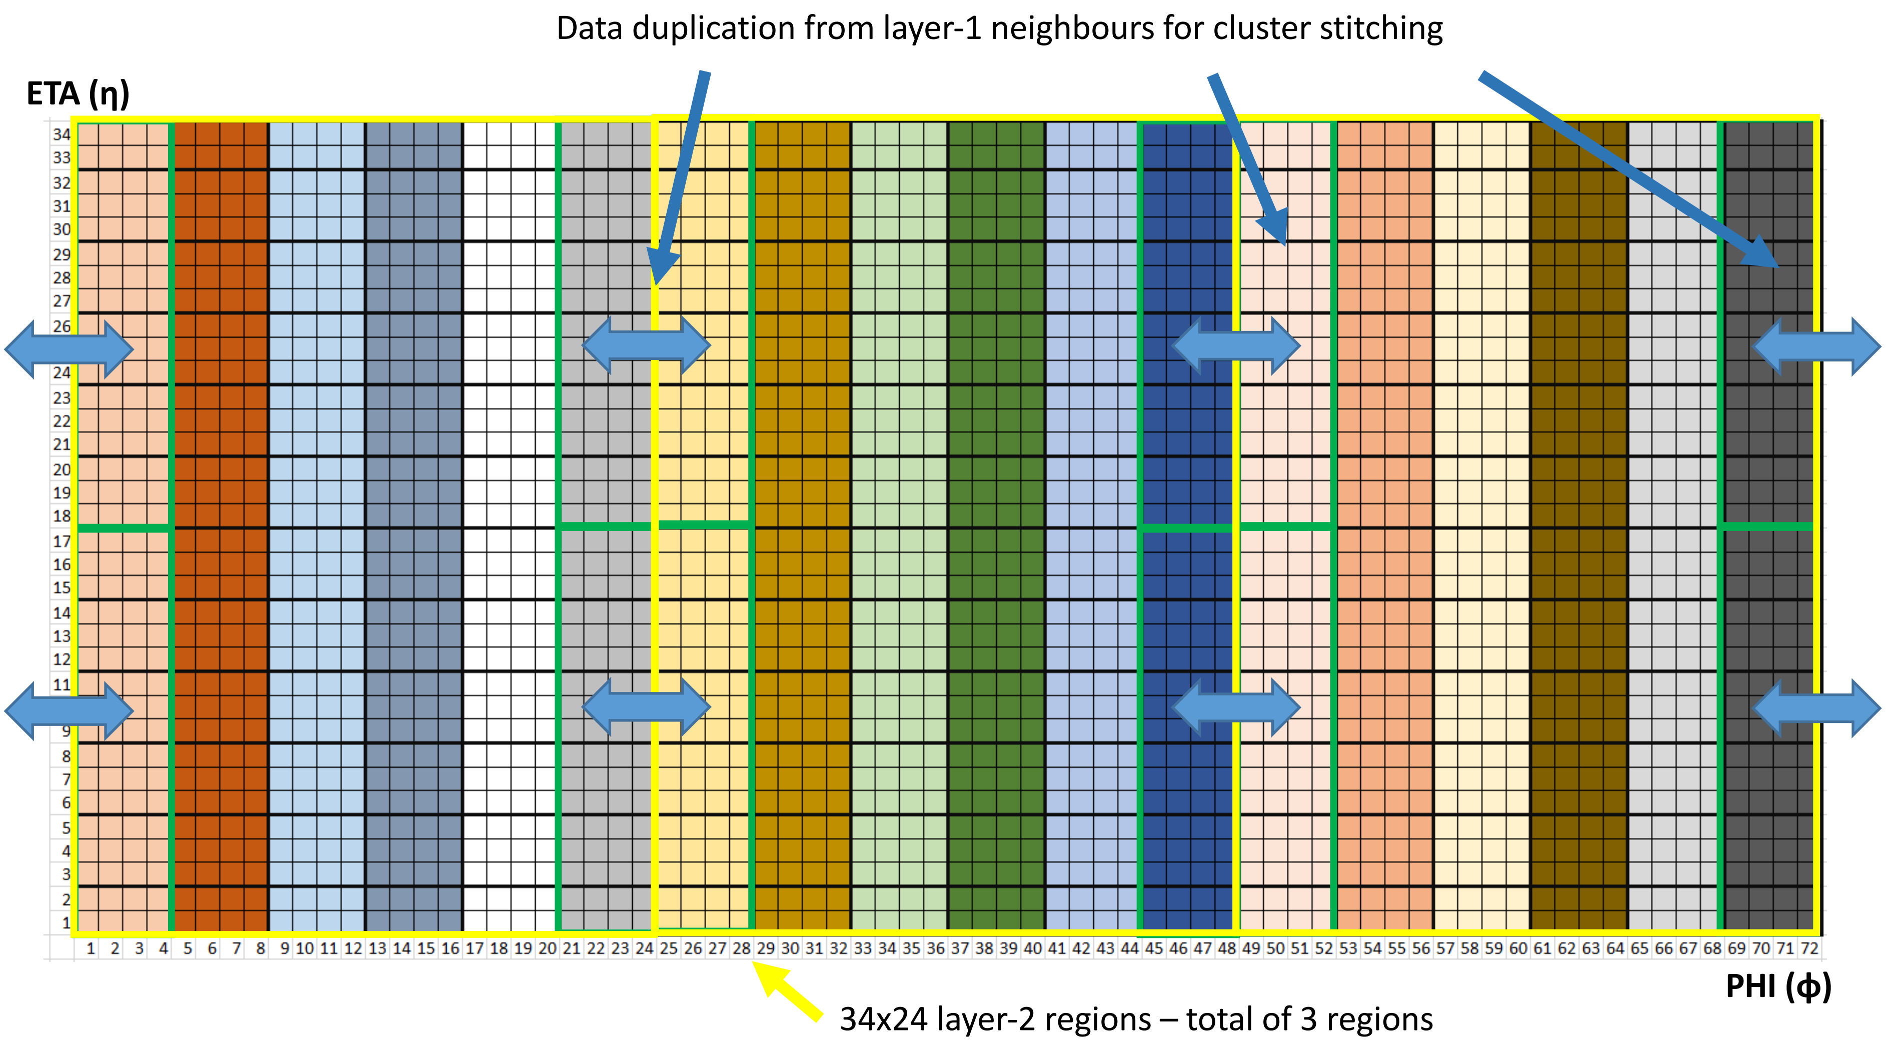

System Regional Layout

Layer-1 recap:
- 5x **3η x 4ϕ** and 1x **2η x 4ϕ** ECAL regions
- 1x **16η x 4ϕ** HCAL regions

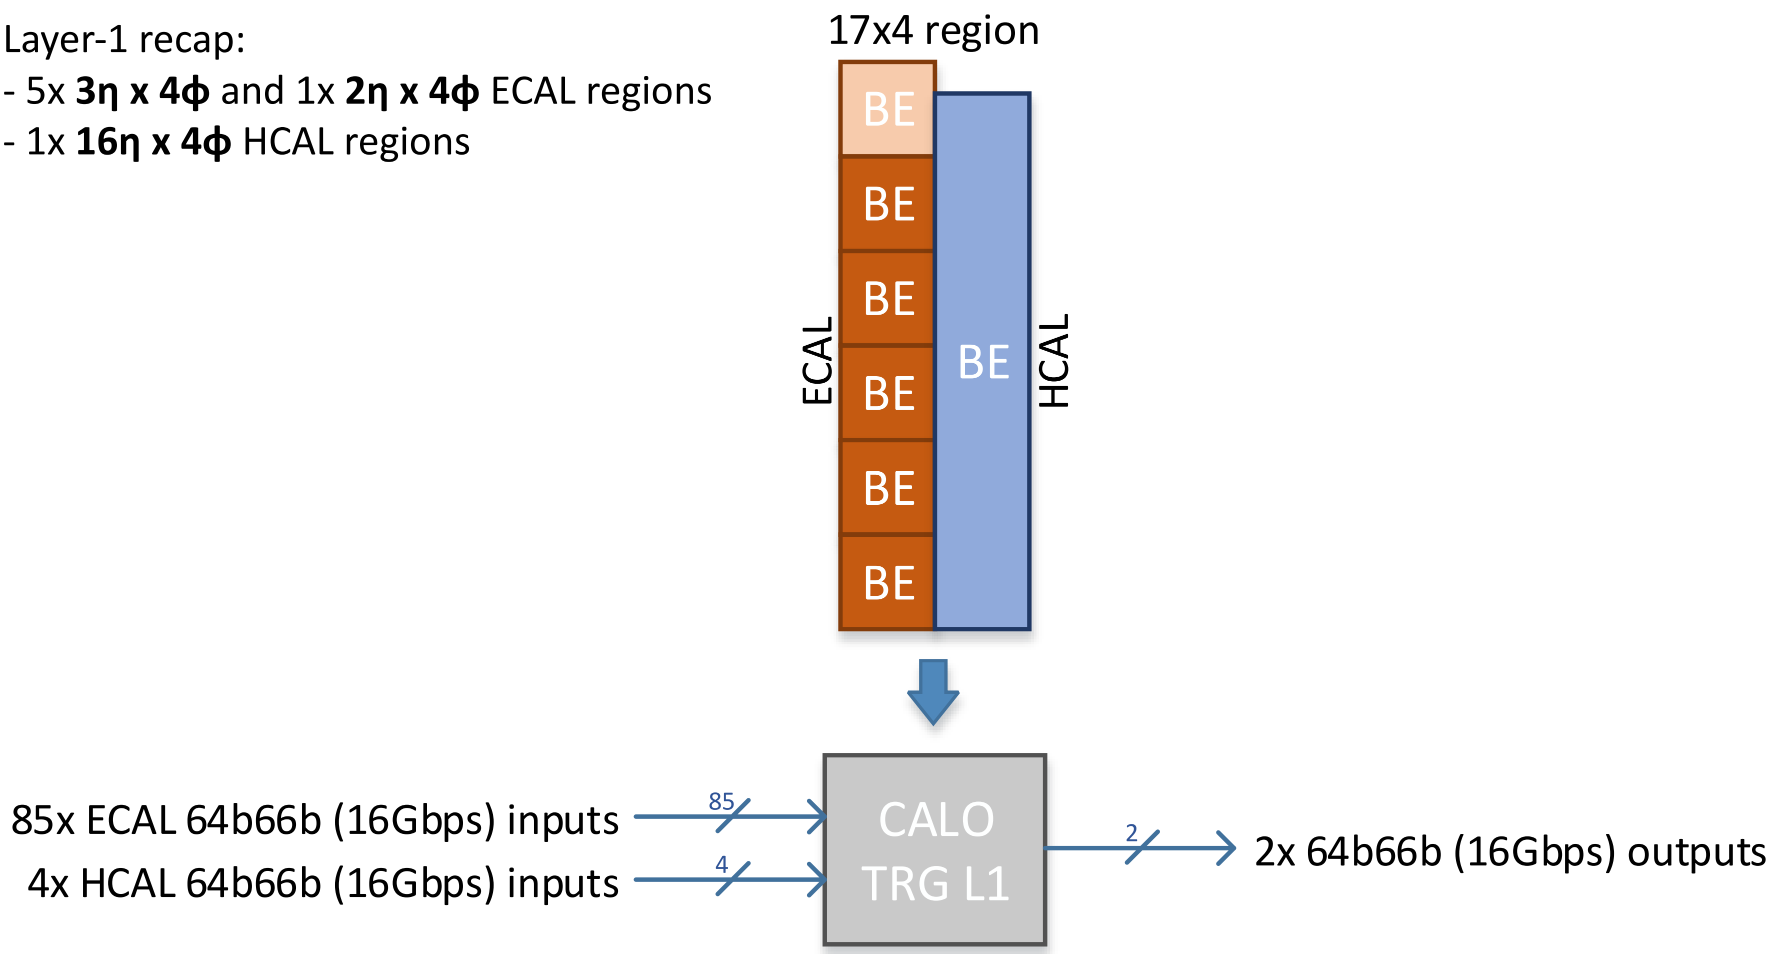

Overall Regional Layout

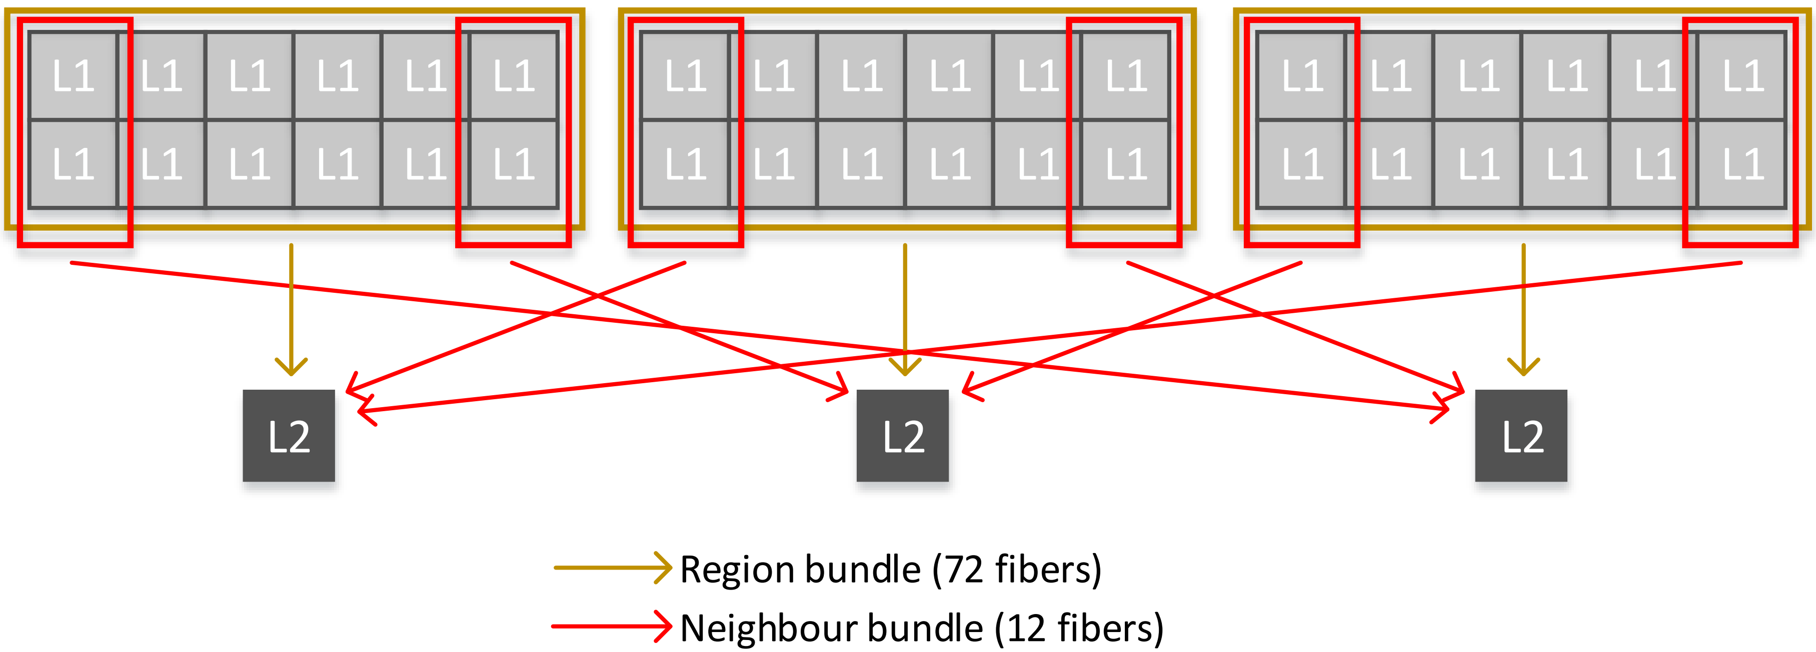

Full System Layout

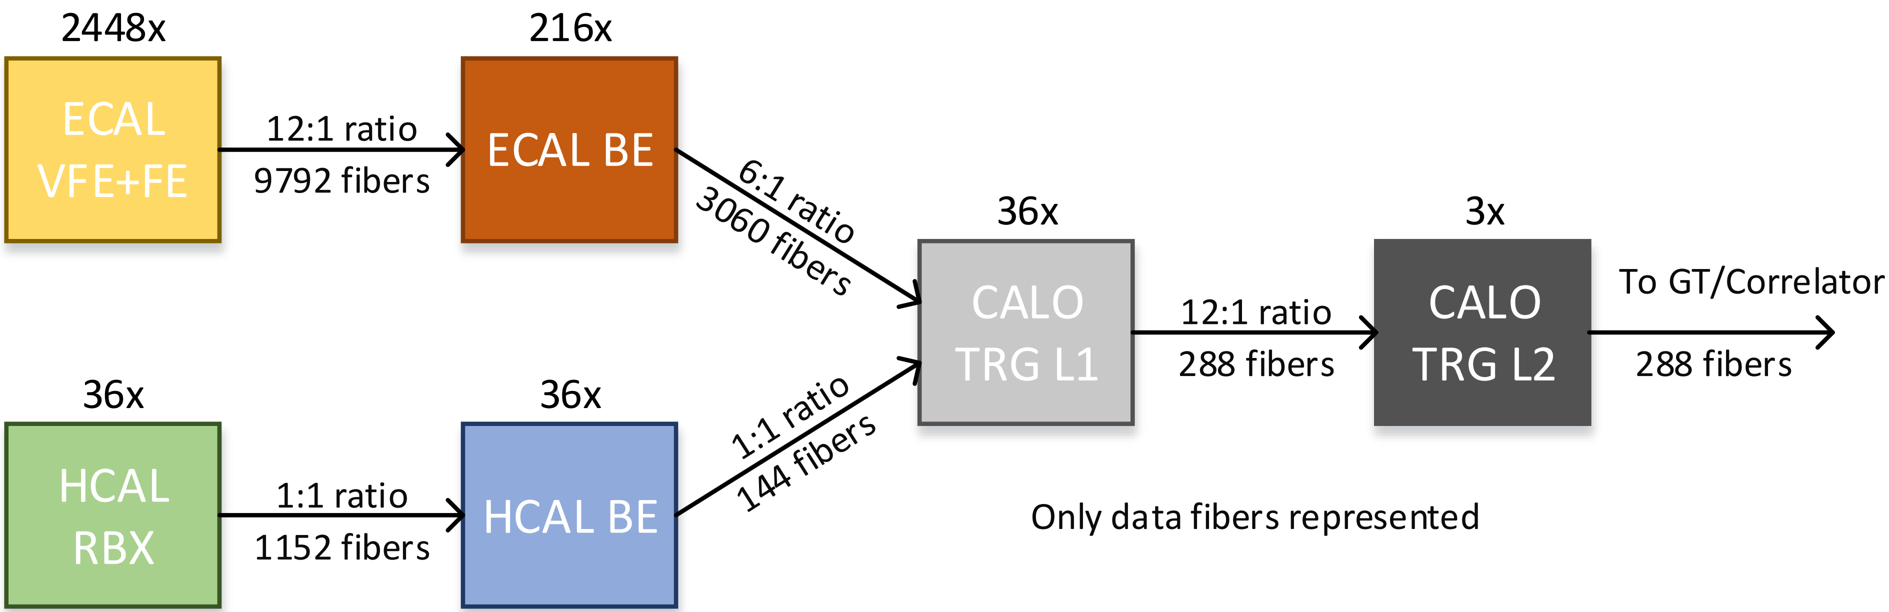

- Ratios reflect **ηxϕ** input regions to output regions
- Counts represent total number of FPGAs per system layer

DAQ readout

- Four 16Gbps lanes are reserved per card for DAQ.
  - Assume that only output data will be readout and at the maximum expected rate of 750 kHz (events per second).
  - Input data can also be readout for test purposes.
  - A total of **64 Gbps per card** is allowed when using four lanes.

- Layer-1:
  - Readout bandwidth for 6 output fibers: **1.8 Gbps per card**.
  - Layer-1 total readout bandwidth (36 cards): **64.8 Gbps**.
- Layer-2:
  - Readout bandwidth for 96 output fibers: **28.8 Gbps per card**.
  - Layer-2 total readout bandwidth (3 cards): **86.4 Gbps**.

R&D Program Starting Point

- U. Wisconsin CTP7 MicroTCA Card for Phase 1 Cal. Trig.
  - 12 MGT MicroTCA backplane links
  - 67 Rx and 48 Tx 10G optical links

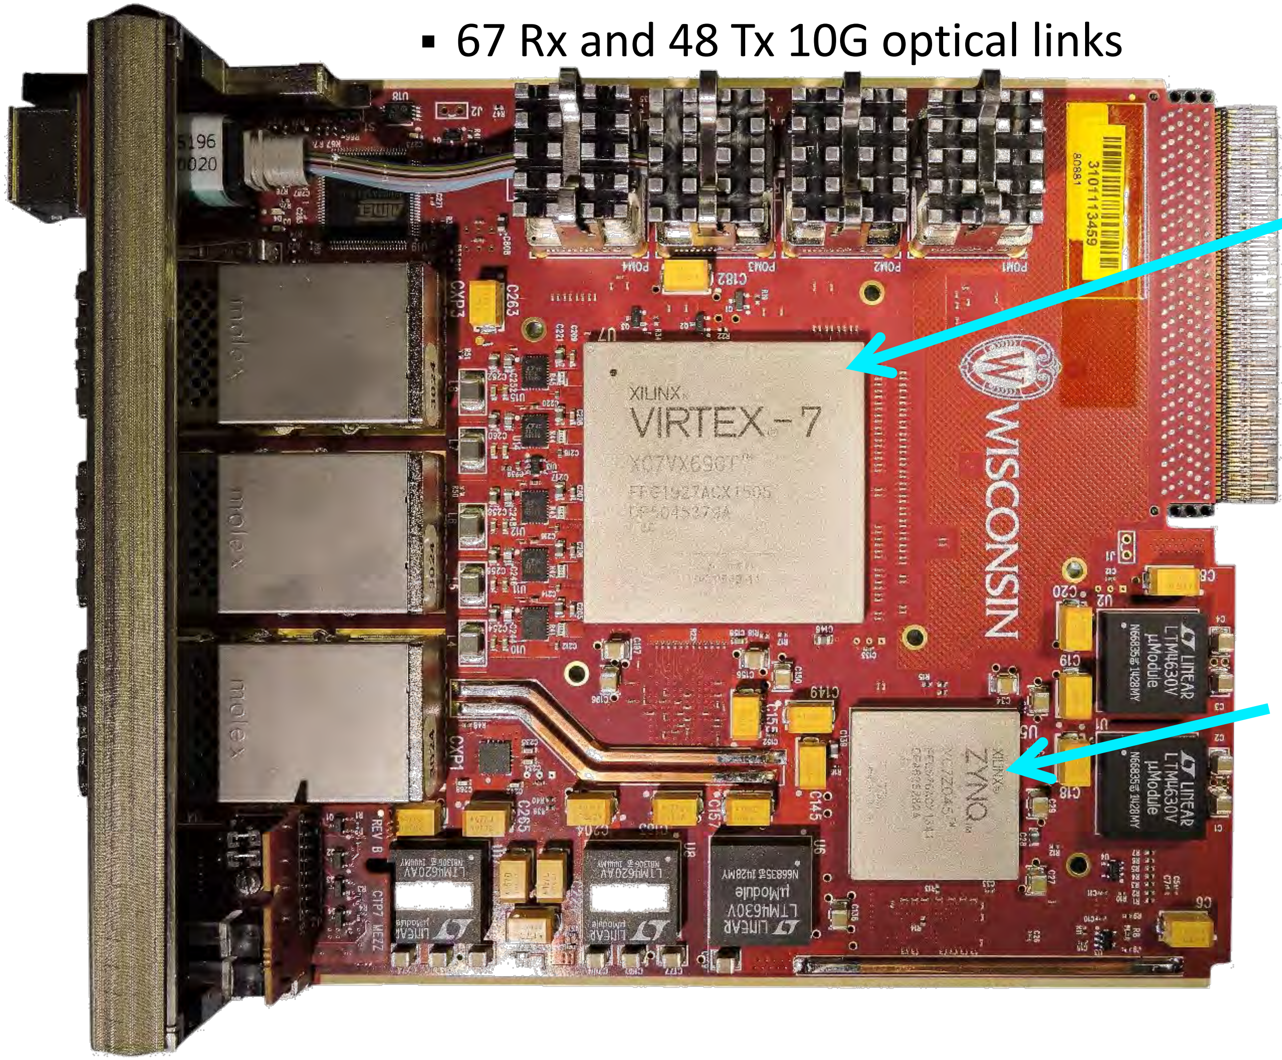

Virtex-7 690T FPGA (Data Processor)

ZYNQ `045 System-on-Chip (SoC) Device (embedded Linux control platform)

CTP7 Deployment: Phase 1 & HL-LHC

- Production:
  - 50 Boards
  - Phase 1 L1 Trigger Deployment:
  - Stage 1 and Stage 2 Layer-1 Calorimeter Trigger
    - 22 CTP7s
  - Stage-1 was main calorimeter trigger for 2015
  - Stage-2 was main Layer-1 calorimeter trigger since 2016
- HL-LHC R&D: Cornell Track Trigger demonstrator test setups
  - 4 CTP7s @ CERN
  - 2nd setup at Cornell: 4 CTP7s
- HL-LHC Cal, Correlator Trigger prototypes: platforms for FW development and testing
- HL-LHC EMU Readout prototype: FW development and testing

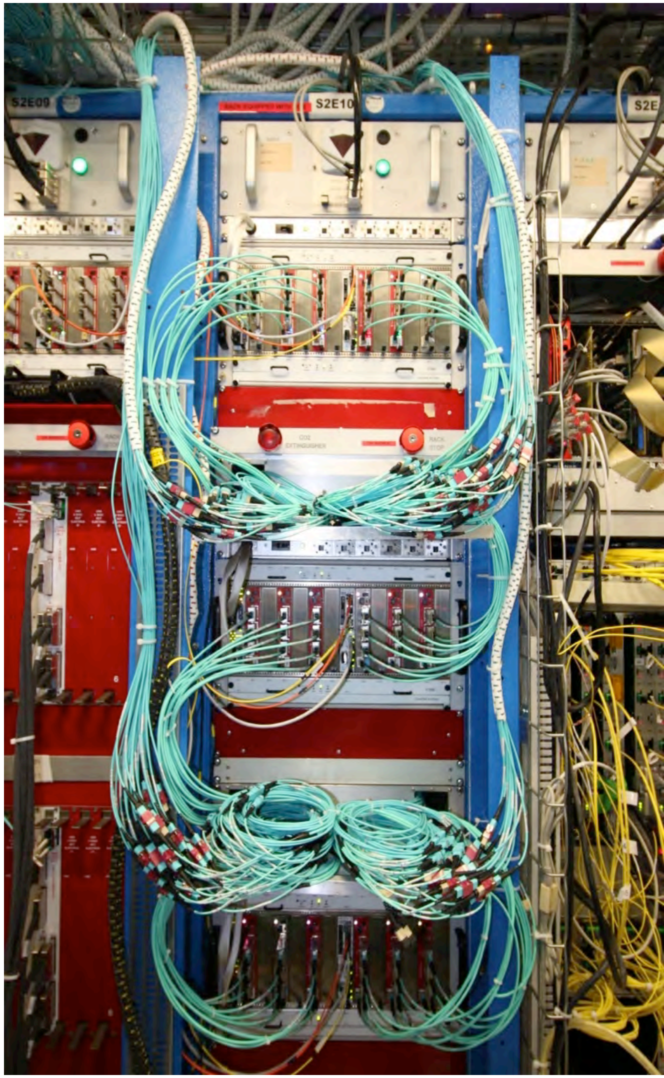

CTP7 Link Integrity in Phase 1

- The Phase 1 Calo L1 CTP7 system has 576 optical inputs from ECAL at 4.8Gbps, 504 HB/HE optical inputs at 6.4Gbps, and 72 HF optical inputs at 6.4Gbps
- CTP7 Integrated Eye Scan capability: non-invasively capture eye diagrams on live operational data upon request
  - Can scan all 1152 input links simultaneously
  - Excellent tool for PM and diagnostic monitoring

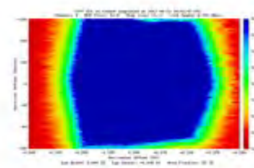

- Automatic Error Handling
  - Packets protected by error-detection codes
  - Payload data is automatically zeroed in firmware for propagation through the trigger algorithms
  - Packets with errors are tagged in the DAQ readout

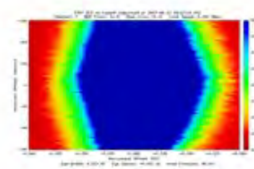

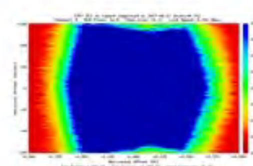

Phase 2 Demonstrator Objectives

1. Explore hardware technologies targeted for the Phase 2 upgrade
   - ATCA Form Factor including Rear Transition Module
   - MGT Link design beyond 10G line rates (16G, 25G)
   - Efficient cooling of next-gen FPGAs
   - Next generation IPMI and embedded Linux solutions
   - Advanced RAM/FPGA interconnections (U. Florida)
2. Identify design blocks suitable for re-use across platforms, either as reference designs or mezzanine boards
3. Provide next-gen platform for ongoing software and firmware R&D work for Phase 2

APd1 Trigger Demonstrator

- General ATCA technology demonstrator, with emphasis on Trigger applications
  - Powerful performance with flexibility
  - Closely related to the ECAL Demonstrator
- Specifications:
  - Single FPGA Design, C2104 Package
  - ~100 Optical Links — Firefly optical modules
    - 14/16G with options to test 25G links as well.
  - Approximately 24 Links to RTM for enhanced versatility
    - RTM includes some of optical links above
  - Embedded Linux and IPMI Controller on Mezzanines for portability and flexibility
  - Deep Memory Mezzanine (U. Florida)

669980062499

APd1 Block Diagram

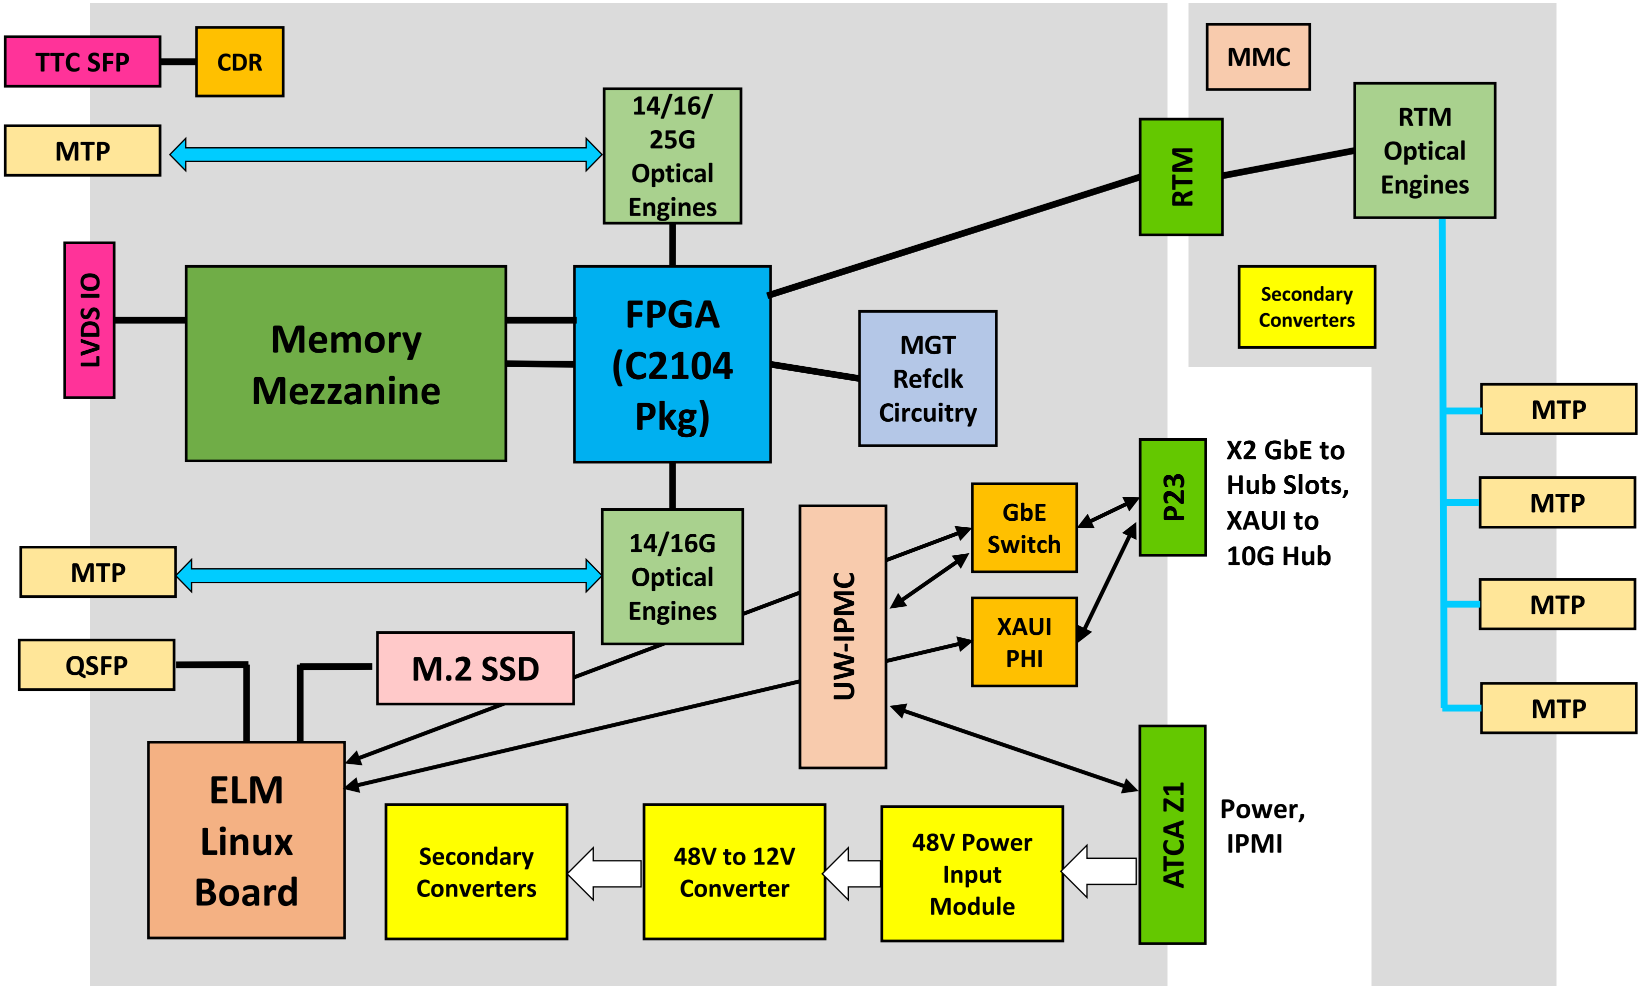

R&D Board Flow

UW-IPMC

- IPMI Carrier Manager host board
- MiniDIMM Form Factor
- ZYNQ ‘020 Based

Embedded Linux Mezzanine (ELM1)

- Embedded Linux Control point
- ZYNQ ‘035-’045 Based
- MGT and FPGA IO to the main board
- 1GbE and 10GbE capable

Controller Development Board (CDB)

- ATCA Blade
- Host development in ATCA crate environment for UW-IPMC and ELM1 boards
- Low-risk proving-ground for mechanical design and ATCA 48V power interface
- No processing FPGA or optical links

ATCA Processor Demonstrator APd1

- ATCA Blade
- Functional demonstrator
- Leverage infrastructure from previous boards in the design flow

IPMI Controller: UW-IPMC

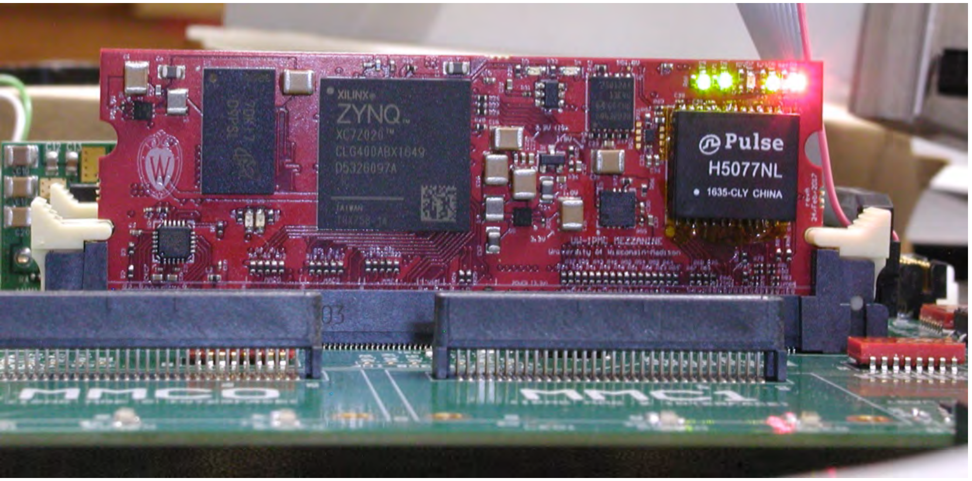

- IPMC: IPMI Controller for ATCA blades
- ZYNQ 7020, RTOS-based application
- I/O Support:
  - Up to 5 MMCs (RTMs, AMCs, etc.) with dedicated IPMB-L (I^2C) bus for each
  - 16 ADC inputs for main board electrical/thermal monitoring and fast fault response
  - 49 3.3V configurable IOs from ZYNQ PL Section
  - 1000BASE-T Ethernet

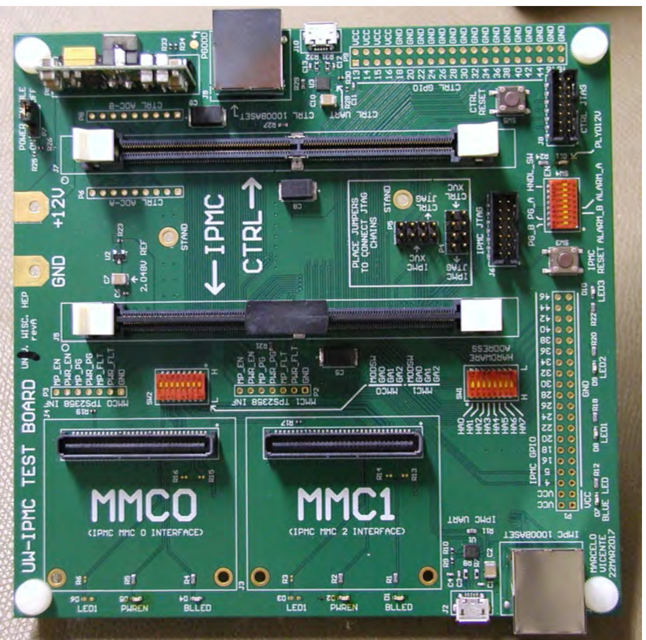

← Top View of Tester Board

ELM1 Embedded Linux Mezzanine

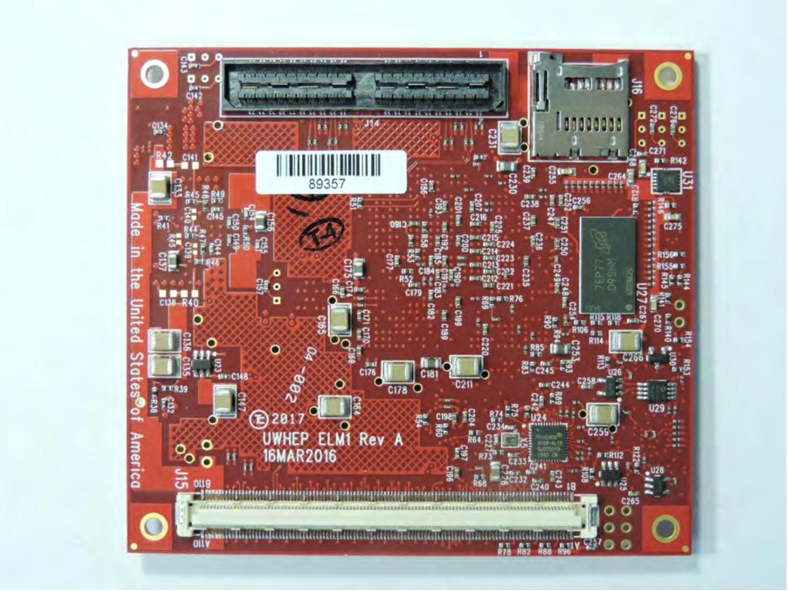

- ZYNQ-based embedded Linux endpoint for ATCA blades
- 84mm × 75mm design, mounts 5mm above main board
- ELM1: gen 1 board with ZYNQ 7000 035/045 device (8 MGT links)
- USB 2.0: 2 ports
- 512 MB of DDR RAM (1066)
- On-board boot sources: QSPI and MicroSD Flash
- Ethernet: GbE and 10GbE capable
- FPGA IO: Over 24 signals @3.3V, 74 high performance signals @1.8V
- Dedicated JTAG Master and Slave ports

Controller Development Board

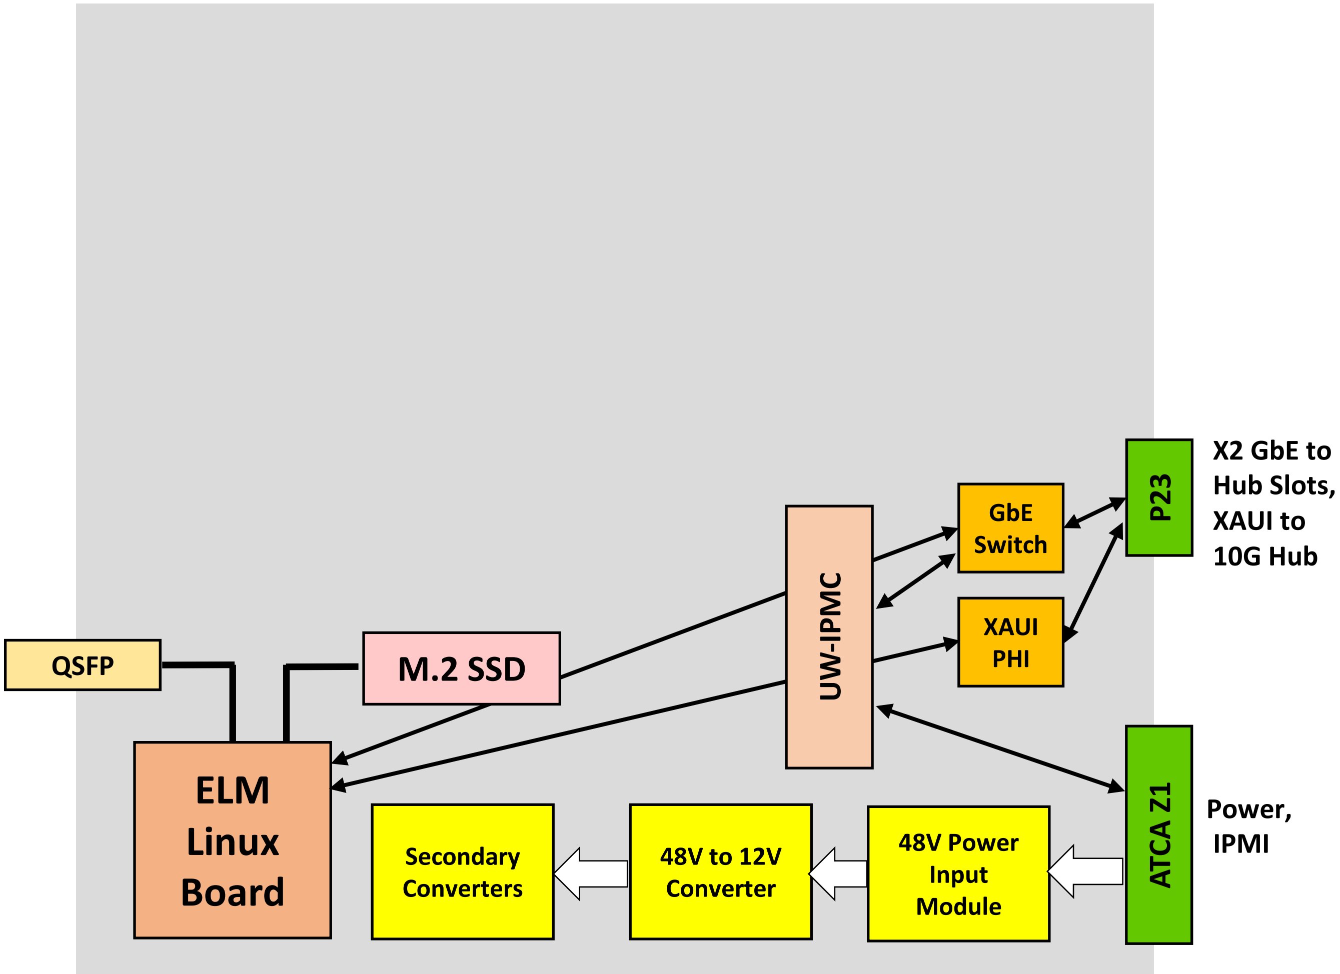

Controller Development Board

- Simple ATCA Board
- Essentially the infrastructure half of the APd1
  - Power
  - IPMI and Embedded Linux connectivity
- Board for verifying mechanical details of ATCA card design, platform for controller development within the ATCA crate
- Allow controller SW/FW development to get out in front of the APd1 hardware design
- *Have Controllers ready for APd1 bring-up!*

2017-2018 R&D Milestones

- 2017 Q2 (30-June-2017): ATCA Control Infrastructure Mezzanines Fabricated
  - UW-IPMC and ELM1 boards fabricated
  - Status: cards under test in the lab
- 2017 Q3 (30-September-2017): ELM1 Standalone Test Board Design Complete
- 2017 Q4 (31-December-2017): CDB Design Complete
  - ATCA test board
- 2018 Q1 (31-March-2018): ATCA Control Infrastructure Demonstrator Assembled
  - CDB with UW-IPMC and ELM1 mezzanine boards
- 2018 Q2 (30-June-2018): ATCA Control Infrastructure Mezzanine First SW/FW release
- 2018 Q3 (30-September-2018): APd1 Produced
- 2018 Q4 (31-December-2018): APd1 Data connectivity test

2019-2020 R&D Milestones

- 2019 Q1 (31-March-2019): APd1 first FPGA firmware infrastructure release
- 2019 Q2 (30-June-2019): UW-IPMC rev.2 design complete
- 2019 Q3 (30-September-2019): ELM2 design complete
- 2019 Q4 (31-December-2019): Subsystem Interconnect test
  - Calorimeter BE → Calorimeter Trigger → Correlator
- 2020 Q1 (31-March-2020): APd2 design complete
- 2020 Q2 (30-June-2020): ATCA Control Infrastructure Mezzanine Second SW/FW release
- 2020 Q3 (30-September-2020): APdx second FPGA firmware infrastructure release
- 2020 Q4 (31-December-2020): Pre-production Complete

Summary

- Barrel Calorimeter Trigger Upgrade meets technical performance requirements
- Trigger Designs are based on similar technologies to Phase-1
- Trigger Upgrade uses common ATCA hardware platform and components also used by other CMS systems
- R&D program starts from successful Phase-1 program
- R&D plan develops the needed infrastructure for control and embedded linux and expedites the demonstrator
- Demonstrator program will complete sufficient testing and validation to launch pre-production.

Backup

FPGA package support

Pkg	mm	Kintex® UltraScale™							Kintex UltraScale+™						Virtex® UltraScale							Virtex UltraScale+					
		KU025	KU035	KU040	KU060	KU085	KU095	KU115	KU3P	KU5P	KU9P	KU11P	KU13P	KU15P	VU065	VU080	VU095	VU125	VU160	VU190	VU440	VU3P	VU5P	VU7P	VU9P	VU11P	VU13P
A784	23		■	■																							
B784	23								■	■																	
A676	27		■	■					■	■																	
B676	27								■	■																	
A900	31		■	■																							
D900	31								■	■		■															
E900	31										■		■														
A1156	35	■	■	■	■		■					■		■													
A1517	40				■	■		■																			
C1517	40						■								■	■	■					■					
D1517	40							■								■	■	■									
E1517	40											■		■													
A1760	42.5													■													
B1760	42.5					■	■	■								■	■	■									
E1760	42.5													■													
D1924	45							■																			
F1924	45					■		■																		■	
A2104	47.5[1]							■								■	■	■					■	■	■		■
B2104	47.5[1]						■	■								■	■	■	■	■			■	■	■	■	■
C2104	47.5[1]																■	■	■	■			■	■	■	■	■
D2104	47.5[1]																								■	■	■
B2377	50																				■						
A2577	52.5																			■					■	■	■
A2892	55																				■						

Legend
■ Device
—— Migration Path

Alternative Architecture Studies

- Example: Use a smaller and less expensive FPGA (below)
  - Fewer links per card → more cards, more layers (latency), awkward geometry, more complexity, more cost
    - More details upon request in parallel session.
- Example: Use two cheaper FPGAs per card
  - Large usage of links and circuitry for data exchange, dividing logic leads to inefficiencies, complex clocking to synchronize, more cost

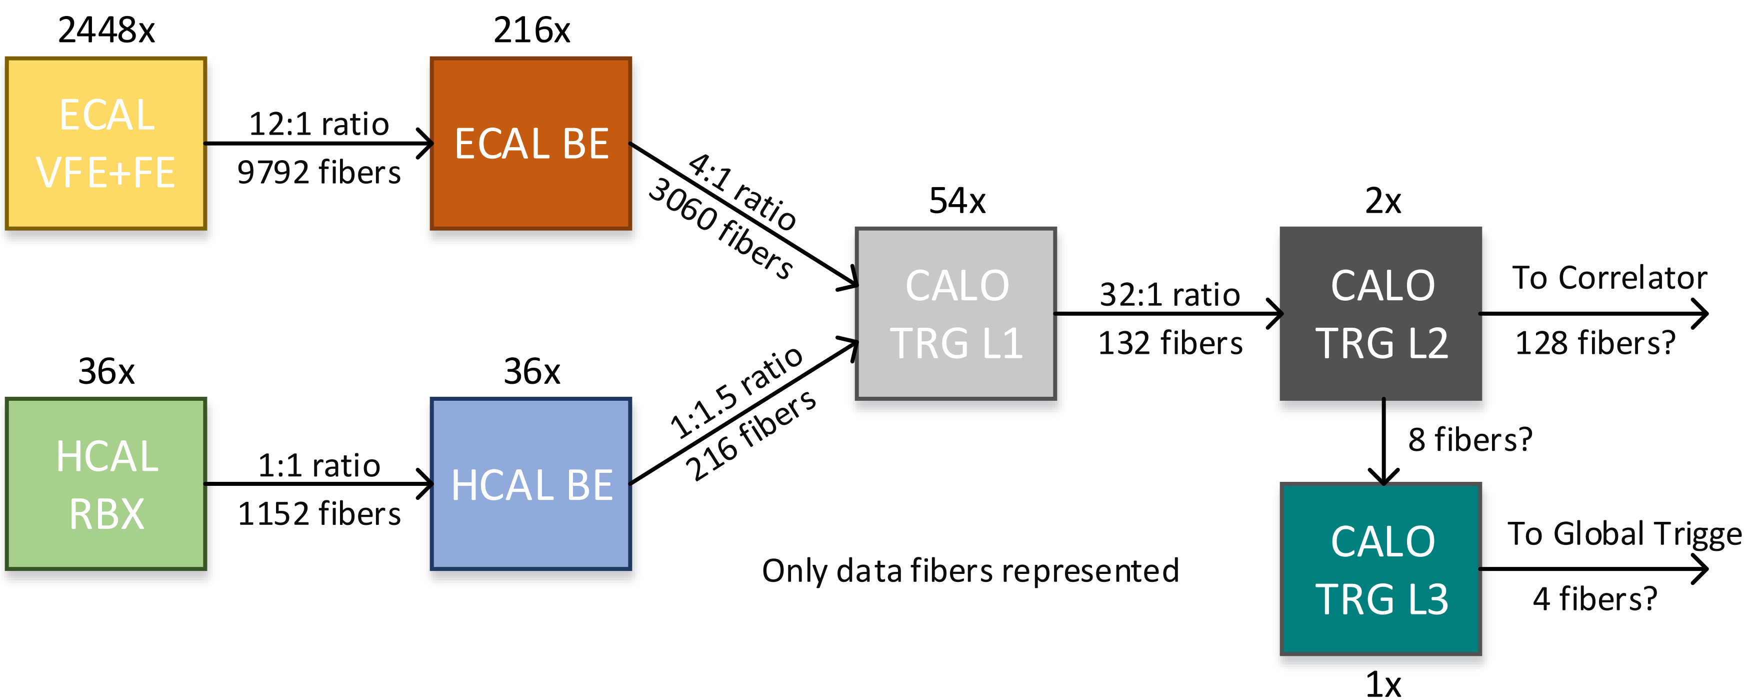

Ratios reflect **ηxϕ** input regions to output regions

Alternate System Layout (2)

- **68 optical links available for data reception and transmission.**

- Layer-1 partitions detector in **11η x 4ϕ** and **12η x 4ϕ** regions – total of 54 regions.
  - ECAL inputs: 4x **3η x 4ϕ** regions (mid-eta region) – 60 ECAL fibers.
  - HCAL inputs: 2x **16η x 4ϕ** region (mid-eta region) – 6 to 8 HCAL fibers.
  - Outputs: 2-4 fibers with regional clusters and metadata (704bits/BX)
- Layer-2 divides the detector into two **ϕ** halves due to input limitations:
  - Receives neighbouring clusters for stitching. Clusters are sent to the correlator.
  - Layer-1 inputs: 27x **11η x 4ϕ or 12η x 4ϕ plus 6 neighbours** – 66 Layer-1 fibers.
  - Outputs: 4 fibers per layer-2 card to send metadata to layer-3 (1408bits/BX) and 64 fibers for each half **ϕ** to send clusters to the correlator.

Alternate System Layout (3)

- A single Layer-3 card will have all metadata from the detector available, allowing the computation of triggerable objects that are then sent to the Global Trigger directly.
  - Layer-2 inputs: Total of 8 fibers with metadata with a clear separation in **ϕ**.
  - Includes the required standalone calorimeter trigger.
  - 68 outputs links available to send trigger objects to the Global Trigger.
  - Could potentially be implemented on a layer-2 card.

- A total of 54 layer-1, 2 layer-2 and 1 layer-3 FPGAs are required: **57 FPGAs.**
  - A total of **140 fibers** are required to send the data between layers.

Alternate System Layout (4)

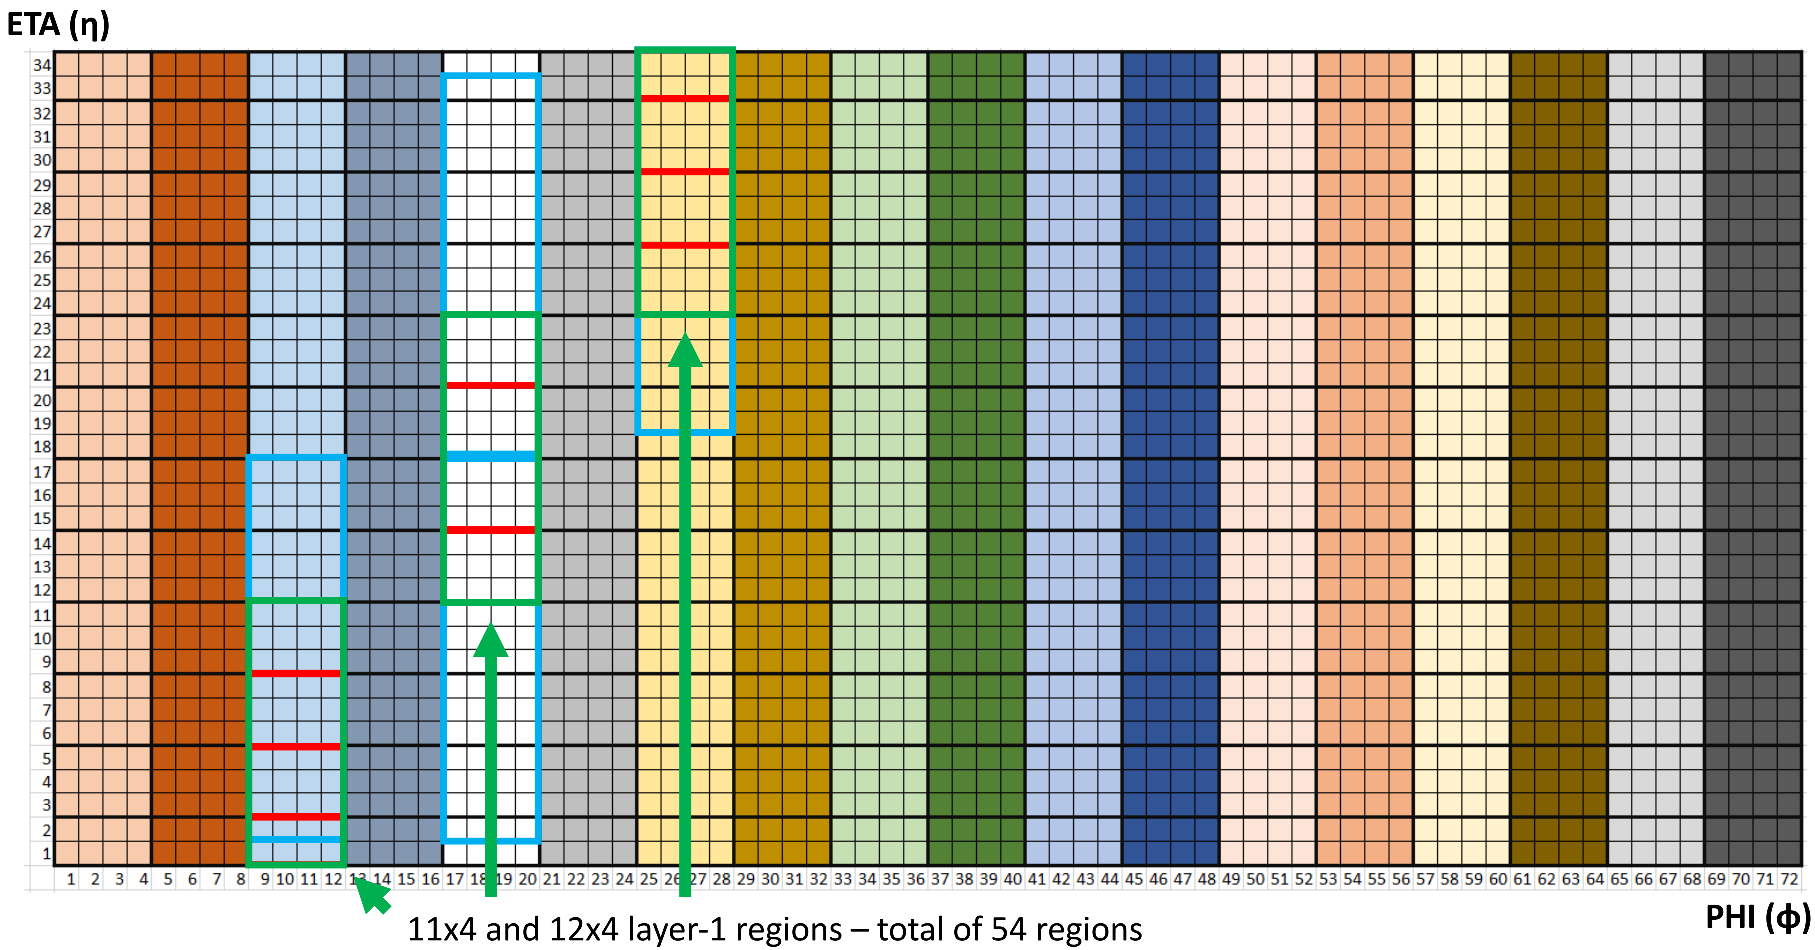

Alternate System Layout (5)

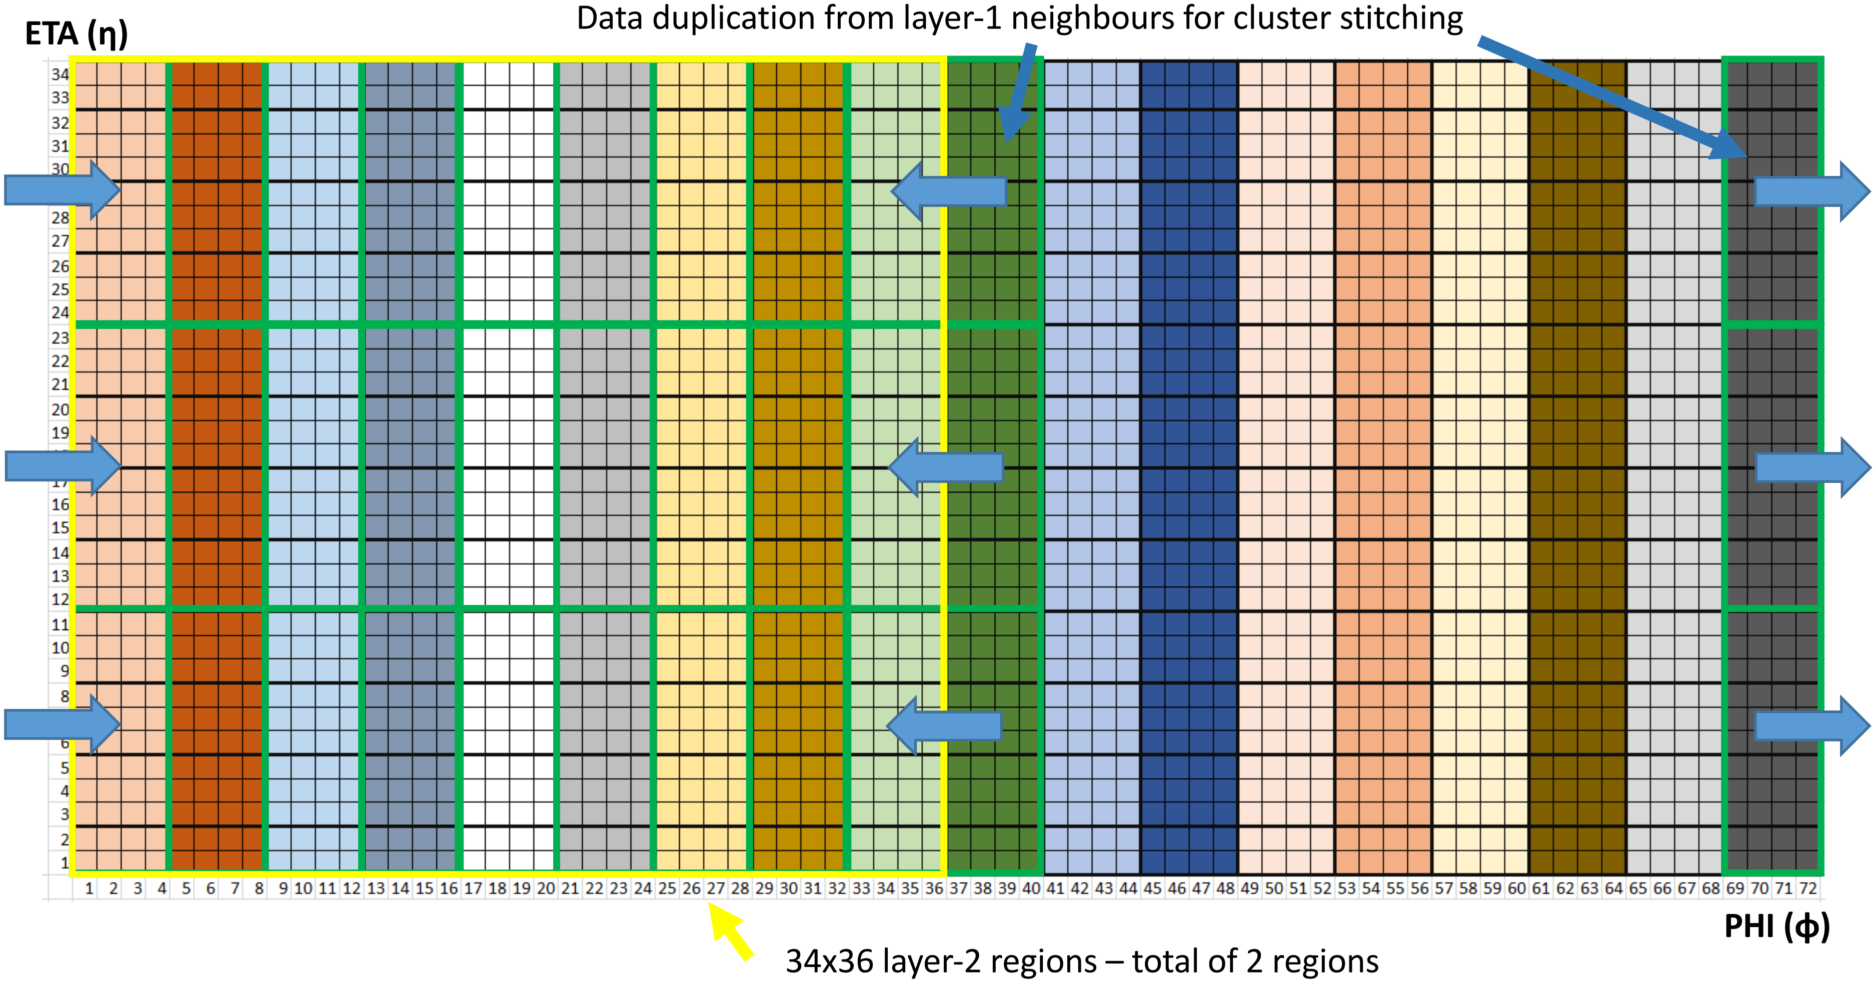

Alternate System Layout (6)

Layer-1 recap for eta edges:
- 3x **3η x 4ϕ** and 1x **2η x 4ϕ** ECAL regions
- 1x <u>partial</u> **16η x 4ϕ** HCAL region

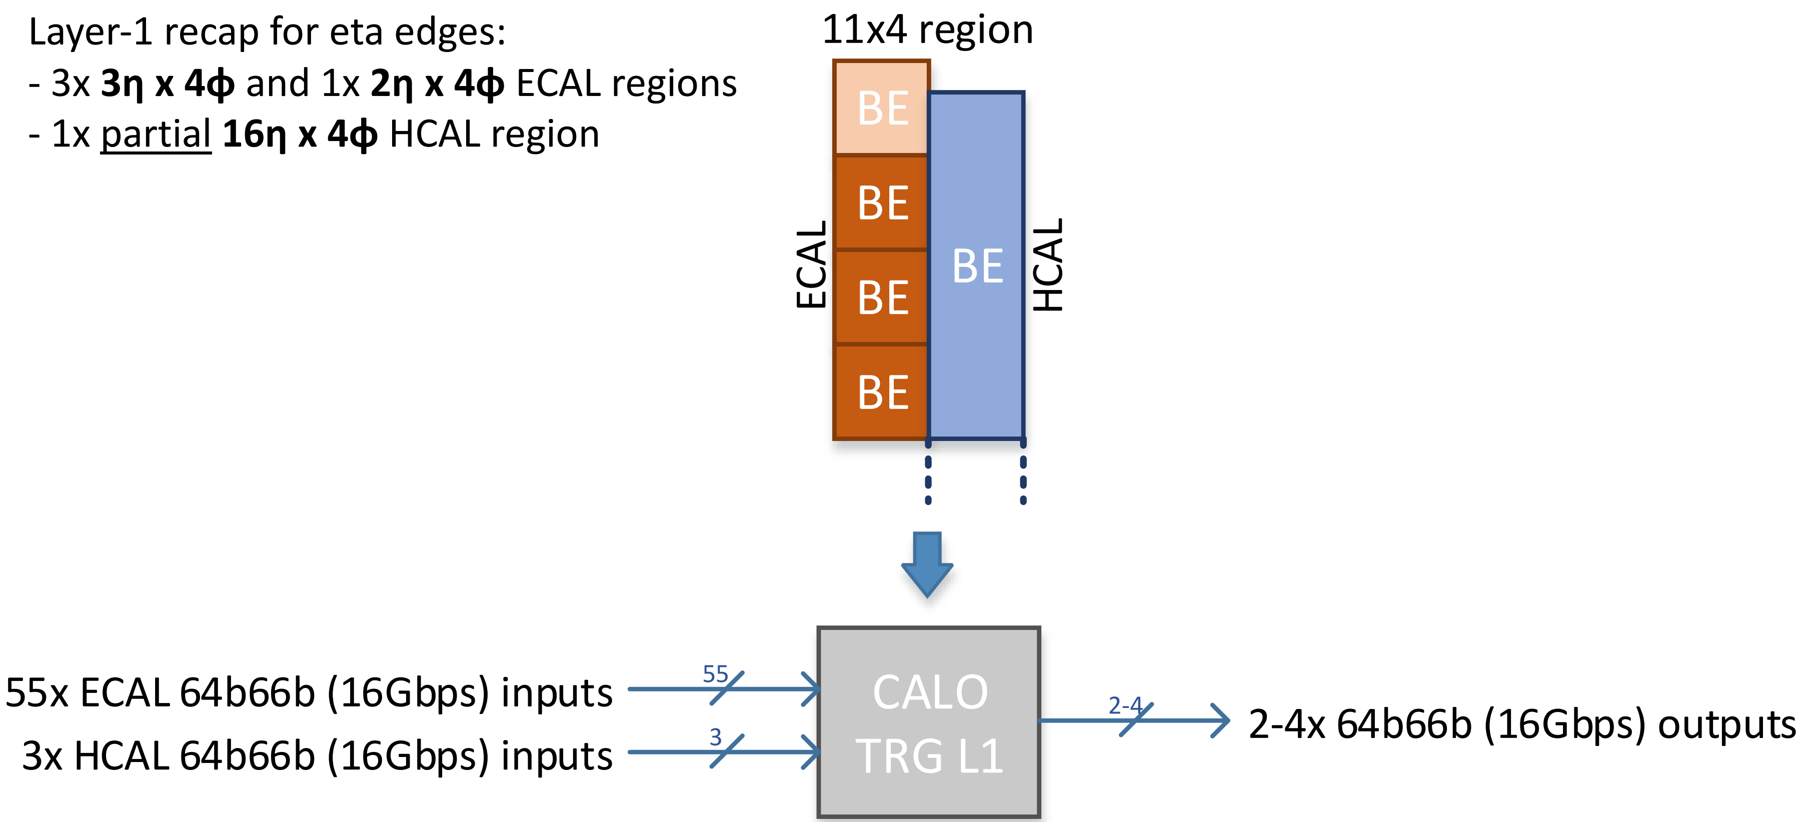

Alternate System Layout (7)

Layer-1 recap for central eta:
- 4x **3η x 4ϕ** ECAL regions
- 2x partial **16η x 4ϕ** HCAL region

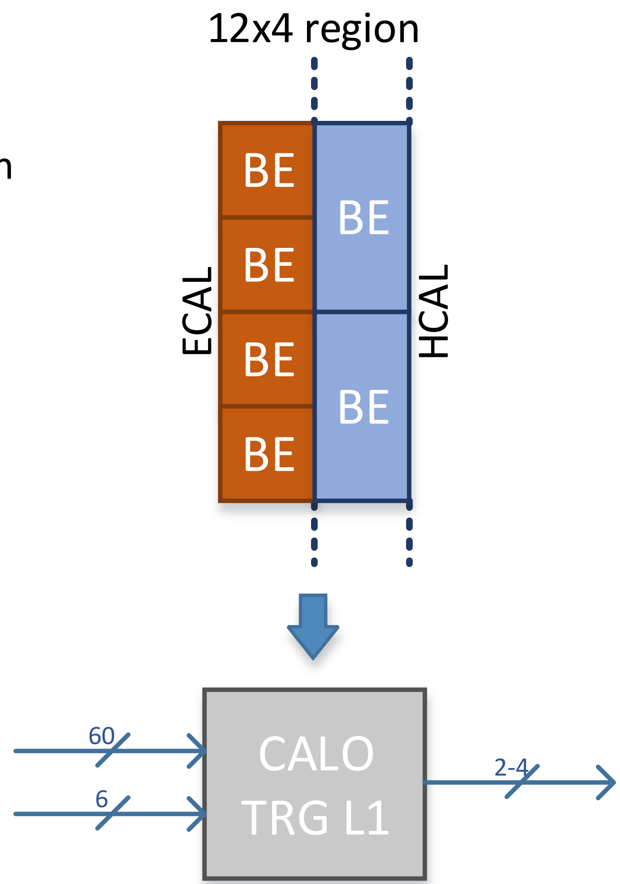

60x ECAL 64b66b (16Gbps) inputs
6x HCAL 64b66b (16Gbps) inputs
2-4x 64b66b (16Gbps) outputs

Alternate System Layout (8)

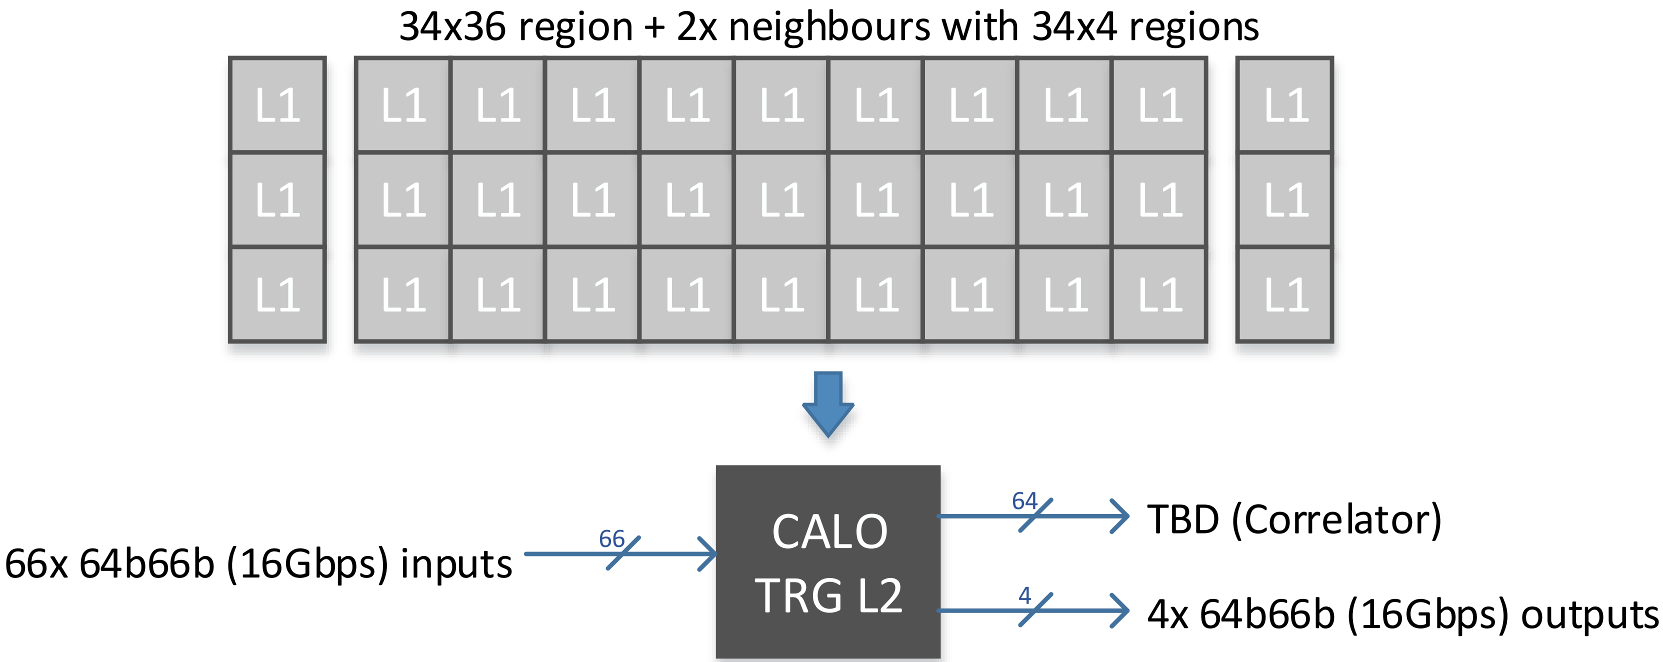

Layer-2 recap:
- 27x **11η x 4ϕ or 12η x 4ϕ** layer-1 regions
- 6x layer-1 neighbouring regions

Alternate System Layout (9)

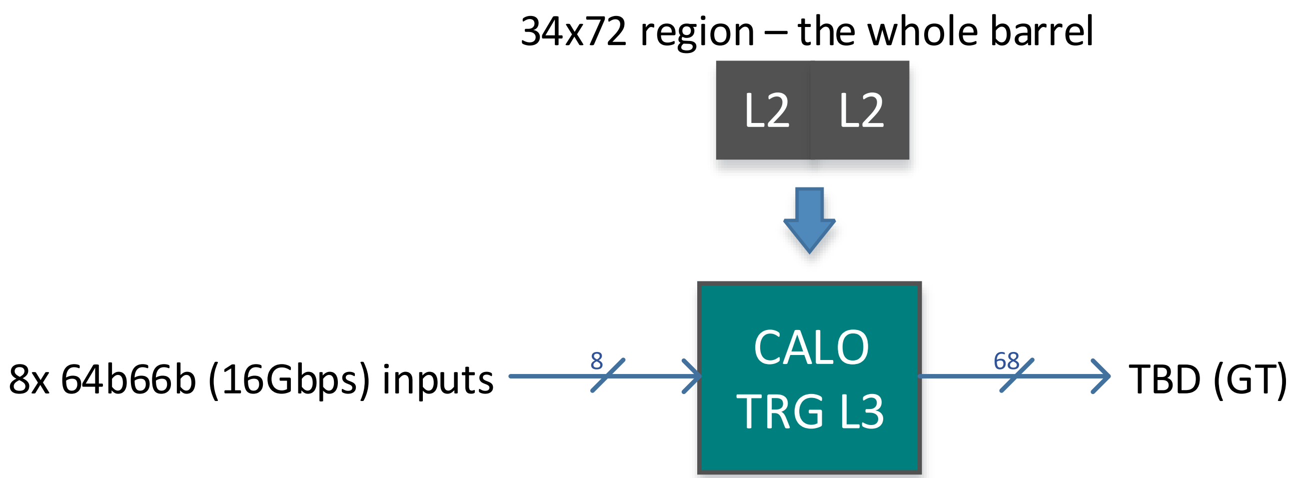

Layer-3 recap:
- 2x **34η x 36ϕ** layer-2 regions